CLAIMS

1. A general purpose computer system having multiple nodes, comprising:

at least one processor executing method acts to promote tolerance of faults in the system, the method acts comprising:

based at least in part on the faults, determining a set of nodes; and

using nodes in the set of nodes only as points on routing paths of messages, and not using any node in the set of nodes for sending or receiving messages.

2. The system of Claim 1, wherein the set is a lamb set.

3. The system of Claim 2, wherein the act of determining undertaken by the processor includes:

finding small sets of partitions of candidate lamb nodes, each partition including a representative node.

4. The system of Claim 3, wherein the act of finding undertaken by the processor includes:

partitioning nodes in the system into maximal intervals of sequential nodes, no node in an interval being a faulty node.

5. The system of Claim 4, wherein the act of finding undertaken by the processor further includes:

returning at least some intervals as at least one set of partitions.

6. The system of Claim 3, wherein the act of determining undertaken by the processor includes:

determining a reachability from at least one representative node to at least another representative node; and

using the reachability to establish a solution set, such that any node in the solution set can reach any other node in the solution set in at most k rounds.

7. The system of Claim 6, wherein the act of determining a reachability undertaken by the processor includes:

computing at least one reachability matrix, using the solution set.

8. The system of Claim 7, wherein the act of determining a reachability undertaken by the processor further includes:

computing at least one intersection matrix.

9. The system of Claim 8, wherein the act of determining a reachability undertaken by the processor further includes:

returning at least one product of at least one reachability matrix and at least one intersection matrix.

10. The system of Claim 6, wherein k equals two.

11. The system of Claim 6, wherein the act of determining undertaken by the processor includes:

minimizing a solution set using at least one weighted graph G.

12. The system of Claim 11, wherein the weighted graph is a weighted bipartite graph.

13. The system of Claim 11, wherein the act of minimizing undertaken by the processor includes:

finding at least one vertex cover C of the graph G.

14. The system of Claim 13, wherein the act of minimizing undertaken by the processor further includes:

using selected elements of the vertex cover C, establishing the lamb set.

15. The system of Claim 1, wherein membership in the set of nodes depends at least partially on a number of processors in a node that are malfunctioning or not functioning.

16. A computer program device comprising:

a computer program storage device readable by a digital processing apparatus; and

a program on the program storage device and including instructions executable by the digital processing apparatus for promoting fault tolerance in a multi-node system, the program comprising:

means for designating a lamb set of nodes in the multi-node system to be used for routing messages within the system.

17. The device of Claim 16, wherein the lamb set of nodes contains nodes that are used only in messages routes.

18. The device of Claim 16, further comprising means for finding small sets of partitions of prospective lamb nodes, each partition including a representative node.

19. The device of Claim 18, wherein the means for finding includes:

means for partitioning nodes in the system into maximal intervals of sequential nodes, no node in an interval being a faulty node.

20. The device of Claim 19, wherein the means for finding includes:

means for returning at least some intervals as at least one set of partitions.

21. The device of Claim 18, wherein the means for designating includes:

means for determining a reachability from at least one representative node to at least another representative node; and

means for using the reachability to establish a solution set, such that any node in the solution set can reach any other node in the solution set in at most k rounds.

22. The device of Claim 21, wherein the means for designating includes:

means for computing at least one reachability matrix;

means for computing at least one intersection matrix; and

means for returning at least one product of at least one reachability matrix and at least one intersection matrix.

23. The device of Claim 21, wherein k equals two.

24. The device of Claim 18, wherein the means for designating includes:

means for minimizing a solution set using at least one weighted graph G.

25. The device of Claim 24, wherein the weighted graph is a weighted bipartite graph.

26. The device of Claim 24, wherein the means for minimizing includes:

means for finding at least one vertex cover C of the graph G.

27. The device of Claim 26, further comprising:

means for using selected elements of the vertex cover C to establish the lamb set.

28. The device of Claim 16, wherein membership in the lamb set of nodes depends at least partially on a number of processors in a node that are malfunctioning or not functioning.

29. A method for promoting fault tolerance in a multi-node system, comprising the acts of:

for each of k rounds, finding multiple partitions of nodes, each partition having a representative node;

for each representative node, determining whether the node can reach at least one predetermined other representative node within a predetermined criteria;

minimizing the number of nodes and/or partitions using a weighted graph to establish a routing set of nodes; and

returning the routing set of nodes for use thereof in routing messages through the system in the presence of one or more node and/or link faults.

30. The method of Claim 29, wherein the number of rounds is at most two.

31. The method of Claim 29, wherein the number of rounds is two and only two.

32. The method of Claim 29, wherein the weighted graph accounts for at least one node weight, the node weight being based at least on a number of operational processors in the node.

33. The method of Claim 29, wherein the routing set of nodes is a lamb set containing nodes that are used only for routing messages.

34. The method of Claim 29, further comprising finding small sets of partitions of prospective lamb nodes, each partition including a representative node.

35. The method of Claim 34, comprising partitioning nodes in the system into maximal intervals of sequential nodes, no node in an interval being a faulty node.

36. The method of Claim 35, comprising returning at least some intervals as at least one set of partitions.

37. The method of Claim 29, wherein the act of determining whether the node can reach at least one predetermined other node comprises:

determining a reachability from at least one representative node to at least another representative node; and

using the reachability to establish the routing set, such that any node in the routing set can reach any other node in the routing set in at most k rounds.

38. The method of Claim 37, wherein the act of determining whether the node can reach at least one predetermined other node comprises:

computing at least one reachability matrix;

computing at least one intersection matrix; and

returning at least one product of at least one reachability matrix and at least one intersection matrix.

39. The method of Claim 29, wherein the act of minimizing includes:

finding at least one vertex cover of the weighted graph; and

using selected elements of the vertex cover to establish the routing set.

40. The method of Claim 29, wherein the weighted graph is a weighted bipartite graph.